Content-based Multi-path Routing in Structured Cyclic Overlays

Muhammad Shafique
School of Computer Science, University of Waterloo
muhammad.shafique@uwaterloo.ca

arXiv:1512.06425v1 [cs.DC] 20 Dec 2015

***Abstract*—Acyclic overlays used in broker-based *Publish/subscribe (PS)* systems provide only one path for content-based routing between publishers and subscribers. This poses serious challenges in handling network conditions like congestion, and link or broker failures. A cyclic overlay may provide multiple paths between publishers and subscribers; however, there is always one subscription tree available for sending a notification to an interested subscriber, which makes content-based dynamic routing a difficult task. This paper introduces a structured cyclic topology that provides multiple paths between publishers and subscribers. The subscription forwarding algorithm exploits the structured nature of the proposed overlay topology and uses a clustering technique to generate shortest-lengths subscription trees without generating duplicate messages. We implemented static and *intra-cluster* dynamic routing algorithms in the proposed overlay topology for our subscription-based PS system, called *Octopi*. Experiments on a cluster testbed show that our approach generates fewer inter-broker messages, and is scalable.**

I. Introduction

Publish/Subscribe (PS) systems are used for many-to-many communication among loosely coupled distributed entities. Data senders (a.k.a. *publishers*) publish the data in form of *notifications*, while receivers (a.k.a. *subscribers*) register their interest in the form of *subscriptions* to receive notifications of interest [10], [2], [7], [21]. Publishers and subscribers are hosted by a set of interconnected brokers which form the *overlay network*. Publishers and subscribers are decoupled in time, space, and flow, and remain anonymous to each other [10], [2]. PS systems is an active area of research due to its increasing popularity and gradual adoption in different application domains [21]. PS systems are the communication substratum in social networking systems [22], [9], business process monitoring [18], software defined networking [20], massive multi-player online games [1], and RSS filtering [21]. PS systems are also used in a number of commercial applications. For example, Yahoo uses PS in PNUTS [6]; Google has introduced Google Cloud PS [13]; Microsoft uses Bing Pulse [12]; LinkedIn uses Kafka [15]; and Wormhole handles messages generated by Facebook users [9].

In a distributed *Content-based Publish/Subscribe (CPS)* system, a *subscription* is a set of *filters* which identify constraints on the contents of notifications [2]. Subscriptions propagate in the overlay network and are saved in the *Routing Table (RT)* of each broker in order to set a subscription tree as a routing path. A notification is routed stepwise from the publisher's host broker to brokers in the routing paths of subscribers with matching subscriptions. This notification routing is called the *content-based* or *filter-based* routing using *Reverse Path Forwarding (RPF)* [4].

Most CPS systems use acyclic topology and offer only one routing path for sending notifications [21]. Routing is made simple in this way; however, a single routing path may have an unpredictable number of brokers without local interested subscribers. This condition causes inefficient use of network resources because a number of brokers in routing paths act as *forwarder-only* brokers and generate *pure forwards* in large overlay networks. Furthermore, applications running on top of a CPS layer require robust messaging when failures happen in the system, or a link congestion is detected [14]. CPS systems that use acyclic overlays have limited flexibility for dealing with varying network conditions. For example, they are not robust in handling situations like load imbalance, link congestion, and broker or link failure. In load balancing, several *coordination messages* have to be exchanged between brokers to select a broker for load shifting [3]. Since subscribers are shifted to new brokers, a number of updates have to be done in RTs as a result of unsubscribe and resubscribe operations. In the event of broker or link failure, the topology repair process is complex due to the importance of maintaining the acyclic property of the overlay. Since brokers are aware of only their direct neighbours at one hop, the topology repair process is a challenging task: no broker has full knowledge of the overlay network, which can result in network partition [14]. Cyclic overlays, on the other hand, can improve performance and throughput by offering multiple paths among publishers and subscribers. Since redundant links are available, the best available path can be selected dynamically to route notifications, reducing delivery delays when a link congestion or failure is detected. Although multiple paths may be available from publishers to subscribers, selecting an optimal routing path is challenging. Each subscription generates one subscription tree at most, which does not allow multi-path routing using the traditional RPF technique. If a link gets congested in a *Subscription-based CPS (SCPS)* system, no alternative path can be used to forward notifications, and dynamic routing is not possible. Content-based routing in cyclic overlays has received little attention and, to the best of our knowledge, no SCPS system offers dynamic multi-path routing.

This paper introduces the first SCPS system which offers multi-path dynamic routing when congestion in a link is detected. We adopt techniques from *Network Product (NP)*

to introduce the use of cluster-based *structured cyclic overlays* in CPS systems. Our subscription algorithm generates subscription-trees of minimum lengths. The contributions of this paper are as follows.

1) Sec. III introduces how Cartesian Product of graphs can be used to formally describe a structured cyclic overlay. Constraints are introduced for the overlay to eliminate notification cycles (or loops).
2) Sec. IV describes a clustering approach and a lightweight cluster bit vector mechanism to eliminate cycles without padding path information in notifications.
3) A subscription forwarding algorithm that generates subscription-trees of minimum lengths in the cluster-based structured cyclic overlay is described in Sec. V.
4) Sec. VI elaborates on our static and dynamic notification routing algorithms. To provide a comparison with the subscription-based variant of the state-of-the-art routing approach, Sec. VI also describes a routing technique based on identification of the subscription tree in SCPS systems.

Background to the research problem and related work is discussed in Sec. II. Sec. VII evaluates and compares the three routing algorithms presented in Sec. VI. The experimental results show that the static and dynamic routing algorithms for the cluster-based structured overlay perform better than identification-based routing algorithm for the unclustered structured cyclic overlay. The appendix provides details of the algorithms for subscription forwarding and notification routing.

II. Background and Related Work

In this section, we discuss the issue of cycles in content-based routing and solutions proposed by researchers. We also explore why, despite of multiple paths between publishers and subscribers, contemporary approaches are unable to support dynamic routing of notifications in SCPS systems.

Content-based routing generates cycles (or loops) in cyclic overlays. Cycles generate extra messages that route in an overlay indefinitely if they are not detected and discarded. The *Subscription Forwarding Process (SPF)* broadcasts a subscription to set a subscription-tree. The subscription tree is used to forward matching notifications to the subscription issuer using RPF technique [2], [4]. Since multiple paths can be available in a cyclic overlay, a broker may receive duplicates of subscriptions. To avoid duplicate subscriptions, SIENA assigns a unique identification to each broker in a cyclic overlay. A latency-based distance-vector algorithm, with the RPF technique, is used for generating paths and sending notifications [2]. Li et al [16] propose an approach to handle cycles in advertisement-based PS systems. Each advertisement tree is uniquely identified by a tree identification *(TID)* and a subscriber bounds to TIDs of all matching advertisements to receive notifications. The matching process executes only once at host brokers of publishers and TIDs are embedded in notifications for routing toward interested subscribers. A variant of [16] for a SCPS system suggests that each broker in an overlay should be assigned a unique identification. We call this *Broker Identification (BID)*. To discard duplicates, each subscription message carries the BID of the host broker of the corresponding subscriber and is saved on overlay brokers. SFP in cyclic overlays is further explained in Fig. 1(a) which shows a cyclic overlay of six brokers deployed across three geographical regions: R_1, R_2, and R_3. Brokers 1 and 4 are in R_1, brokers 2 and 5 are in R_2, and brokers 3 and 6 are in R_3. The subscribers *S1* and *S2*, connected to *Broker 4* and *Broker 3* respectively, are interested in notifications from a publisher *P* hosted at *Broker 6*. In this paper, an overlay link is represented as $l\langle source, destination\rangle$, where the *source* identifies message sending broker and the *destination* represents the message receiving broker. Broker 4 receives a subscription of S2 (solid arrows) from Broker 5. The link $l\langle 5,4\rangle$ is added in the subscription tree of S2, assuming that Broker 4 discards a second copy of the subscription of S2 received from Broker 1. Similarly, Broker 5 discards a second copy of the subscription of S1 (dotted arrows) received from Broker 2, assuming that the first copy was already received from Broker 4 (there is, therefore, no subscription tree link from Broker 2 to Broker 5). This process indicates that, despite using BIDs with subscriptions, redundant messages are generated to detect and discard duplicates. Although the subscription trees of S1 and S2 in Fig. 1(a) have shortest lengths (number of hops), the approach described in [16] and used for subscription (advertisement) broadcast does not guarantee subscription trees of shortest-lengths. Inter-broker messaging delays are the reason of this limitation. For example, Fig. 1(b) shows a scenario in which links $l\langle 2,5\rangle$ and $l\langle 3,6\rangle$ have heavy network traffic (thick lines) when S2 issues a subscription. Due to a higher delay on $l\langle 2,5\rangle$ and $l\langle 3,6\rangle$, Broker 6 receives the first copy of the subscription of S2 from Broker 5 and discards the second copy when received from Broker 3. Although the host brokers of P and S2 are in the same region (i.e., R_3), S2 receives notifications after they are processed and forwarded by brokers in R_1 and R_2. Subscription trees generated in [2] can be improved by periodically sending subscription messages to find links with the minimum latency; however, this refinement generates extra traffic in the network and is infeasible for a large setup. We argue that subscription-trees with longer lengths result in an increased in-broker computation, extra *inter-broker messages (IMs)*, inefficient use of network bandwidth, and high latency in notification delivery. Ideally, a subscription tree should always has the shortest length, even if some links have high loads when the subscription is issued (e.g., the subscription trees in Fig. 1(c)).

The issue of cycles is also relevant for delivering notifications. For example, the matching process executed at Broker 6 in Fig. 1(a) indicates that a notification n from P should be forwarded onto $l\langle 6,3\rangle$ and $l\langle 6,5\rangle$. Broker 6 creates two copies of n, n_1 and n_2, to forward to Broker 5 and Broker 3 respectively. S1 receives n_1 from Broker 4 and S2 receives n_2 from Broker 3. Unfortunately, the matching process at Broker 3 indicates that the subscription of S1 matches n_2, and a copy of n_2, say $n_2^{'}$, should be forwarded to Broker 2. $n_2^{'}$ ultimately

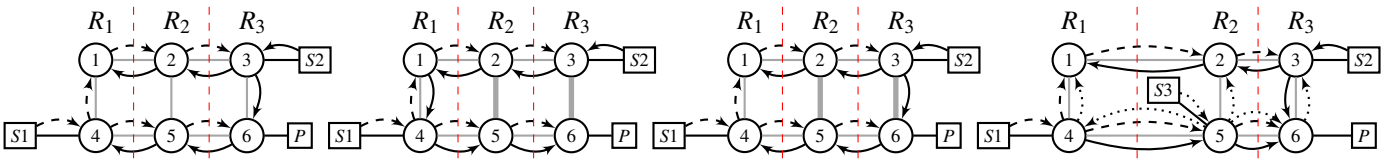

(a) Subscription forwarding in a cyclic overlay when load on the links is normal.

(b) Subscription forwarding when the links $l\langle 2,5\rangle$ and $l\langle 3,6\rangle$ have heavy loaded.

(c) The subscription trees of shortest-lengths when the links $l\langle 2,5\rangle$ and $l\langle 3,6\rangle$ have heavy load.

(d) The BID-based notification routing in cyclic overlays. BIDs assgined to subscriptions are embedded in matching notifications.

Fig. 1: *A cyclic overlay of six brokers. The grey lines indicates links that connect the overlay brokers. The dashed, solid and dotted arrow messages indicate the subscription trees of S1, S2, and S3 respectively. The dashed red lines seprate regions.*

reaches S1 after being forwarded by Broker 4. Again, the matching process at Broker 4 identifies that S2 should also receive a copy of $n_2^{'}$, and this process continues indefinitely until it is identified and stopped. Similarly, n_1, forwarded by Broker 6 onto $l\langle 6,5\rangle$, is received more than once. Adopted from notification routing described in [16], the host broker of a publisher adds BIDs of interested subscribers, when a notification is issued. The RPF technique is used to route notifications along the paths identified by the BIDs. Fig. 1(d) shows subscribers S1, S2, and S3 interested in notifications from P. For any notification m from P, Broker 6 creates two copies, m_1 and m_2, where m_1 with the BID of Broker 3 is forwarded onto link $l\langle 6,3\rangle$, and m_2 with the BIDs of Broker 4 and Broker 5 is forwarded onto link $l\langle 6,5\rangle$. After receiving m_2, Broker 5 forwards one copy of m_2 to S3, and removes the BID of Broker 5 from m_2 before forwarding it to Broker 4. Broker 4 forwards a copy of m_2 to S1 and removes the BID of Broker 4. Since no BID is left in m_2, Broker 4 does not forward m_2 any further. A limitation of the BID-based routing approach is the fact that adding BIDs assigned to matching subscriptions increase the payload of a notification. In a large overlay network, a notification may be received by a high number of subscribers hosted by many brokers and many BIDs may have to be embedded in a notification [12]. In scenarios where scalability is a major requirement, the need to embed many BIDs creates a bottleneck. Octopi uses a cluster-based structured cyclic overlay, which does not require embedding BIDs in notifications.

Another issue with SCPS systems that use cyclic overlays is that there is one subscription tree at most, which is used to forward notifications using the RPF technique. Alternative paths are not available which makes dynamic routing a complex task. Li et al [16] presents a technique for dynamic notification routing in advertisement-based CPS systems. The technique uses intersecting advertisement-trees for a subscriber to offer dynamic routing. The approach does not work when intersecting advertisement trees are not available. Furthermore, a large number of redundant messages are generated in advertisement broadcast and a subscription may has to be delivered multiple times to the same broker [16]. Alternative or intersecting paths are not available in SCPS systems, which makes dynamic routing a complex task, and to the best of out knowledge, is never done before. Octopi uses a cluster-based structured cyclic overlay for dynamic routing. In particular, it offers *inter-cluster* dynamic routing of notifications (cf. Sec. VI). Our two-step subscription forwarding algorithm generates subscription trees of shortest lengths (cf. Sec. V).

III. Structured Cyclic Overlay

Structured cyclic topologies have been used effectively in Product Networks (PN) for many years [5]. A PN is an interconnected network obtained by taking the product of two or more graphs. Many PNs have been proposed as popular topologies, including product tree, tori, hypercube, and butterfly, etc [5]. These topologies have regular structures that make a formal study of properties like symmetry, diameter, fault tolerance, degree, network partition, and possible parallel paths possible. Parallel paths can be used as alternative routes when a node or a link failure occurs. A study shows that the repetition of patterns (or use of template) makes design and analysis of computer networks simple. The use of graph product is effective in describing the design templates that are used to define path redundancy and estimate fault tolerance [19]. We use the Cartesian Product of Undirected Graphs (CPUG) to design large, structured overlay networks based on small graph patterns. In addition, we introduce a clustering technique that divides a structured overlay into identifiable groups of brokers, in order to eliminate cycles without carrying BIDs with a notification.

Cartesian Product of Undirected Graphs — A graph is an ordered pair $G = (V_G, E_G)$, where V_G is a finite set of vertices and E_G is set of edges or links that connect two vertices in G. The number of vertices of G (called *order*) is $\mid G \mid$ (or $|V_G|$). Similarly, the number of edges in G is $\parallel G \parallel$ (or $|E_G|$). Two vertices $u, v \in V_G$ are adjacent or neighbours if the edge $uv \in E_G$. The set of neighbours of a vertex $v \in G$, denoted as $N(v)$, is called the *degree* of v. A graph in which each pair of vertices are connected by an edge is called a *complete graph*. The path distance d between two nodes a, b in a graph is measured in the number of edges between them, and is given as $d\langle a,b\rangle$. The diameter of a graph G, represented as *diam(G)*, is the shortest path between the two most distant nodes in G. A graph product is a binary operation that takes two small graph operands—for example $G(V_G, E_G)$ and $H(V_H, E_H)$—to produces a large graph whose vertex set is given by $V_G \mathrm{X} V_H$. Many types of graph products exist, but the three fundamental types are the Cartesian product, the Direct product, and the Strong product [11]. All three of these graph products have

been investigated extensively in graph theory and used widely in PNs. The core concept behind these graphs products is the rule-based interconnection of vertices of the graph operands. The Cartesian product provides the most suitable link pattern for our research. The link pattern generated by Direct product makes routing complex, and although Strong product provides robust adjacencies (i.e., more edges) between vertices of the operand graphs, a high node degree exerts more load on an overlay broker. Furthermore, Strong product introduces a high level redundancy in overlay links, which makes message routing a difficult task. We discuss CPUG and its applications in the following. Details on the other two types of graph products are available in [11].

The $CPUG$ of two graphs $G(V_G, E_G)$ and $H(V_H, E_H)$ is denoted by $G\square H$, with vertex set $V_{G\square H}$ and set of edges $E_{G\square H}$. Two vertices $(g,h) \in V_{G\square H}$ and $(g',h') \in V_{G\square H}$ are adjacent if $g = g'$ and $hh' \in E_{G\square H}$ or $gg' \in E_{G\square H}$ and $h = h'$. Formally, the sets of vertices and edges of a CPUG are given as.

$$V_{G\square H} = \{(g,h) | g \in V_G \wedge h \in V_H\} \tag{1}$$

$$E_{G\square H} = \left\{\langle (g,h)(g',h')\rangle | \begin{array}{l}(g = g', hh' \in E_H) \\ \vee (gg' \in E_G, h = h')\}\end{array}\right\} \tag{2}$$

The operand graphs G and H are called factors of $G\square H$. CPUG is commutative—that is, $G\square H = H\square G$. Although CPUG of n number of graphs is possible, we are concerned with CPUG of only two graphs in this paper.

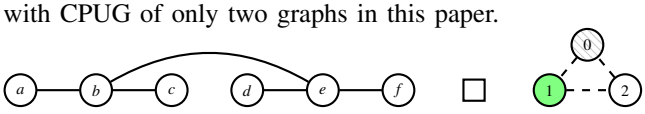

Fig. 2: *Operands of CPUG: Left of $\square$ is G_{af} which is an H-graph; right of $\square$ is G_{cf} which is a triangle.*

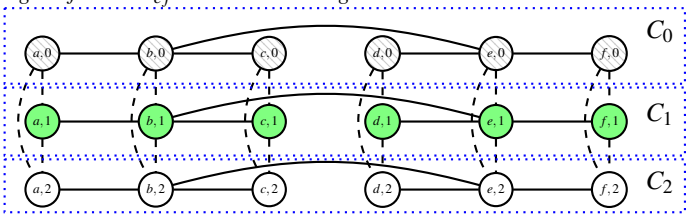

Fig. 3: *Structured Cyclic Overlay Topology (SCOT) generated by CPUG operands shown in Fig. 2.*

Structured Cyclic Overlay Topology — *Structured Cyclic Overlay Topology (SCOT)*, represented as $\mathbb{S}$, is a $CPUG$ of two graphs. One graph, represented by G_{af}, is called *SCOT acyclic factor*, while the second graph operand, represented by G_{cf}, is called *SCOT connectivity factor*. A SCOT has two important properties: (i) *Acyclic Property* emphasizes that the G_{af} must be an acyclic graph, and (ii) *Connectivity Property* requires that G_{cf} must be a complete graph. These properties augment a SCOT with essential characteristics that are used for generating subscriptions trees of shortest lengths. V_{af} and V_{cf} are the sets of vertices of G_{af} and G_{cf}, while E_{af} and E_{cf} are the sets of edges of G_{af} and G_{cf} respectively. For a *singleton graph* of vertex set $\{h\} \subset V_{cf}$, the graph G^h_{af} generated by $G_{af}\square\{h\}$ is called a $G^h_{af} - fiber$ with *index h*. Similarly, for a singleton graph of vertex set $\{m\} \subset G_{af}$, the graph G^m_{cf} generated by $\{m\}\square G_{cf}$ is called a $G^m_{cf} - fiber$ with *index m*. We describe the importance of using indexes in SCOT fibers in Section IV. The definitions of the fibers indicate that, for each vertex of G_{cf}, $CPUG$ generates one replica of G_{af}, and for each vertex of G_{af}, $CPUG$ generates one replica of G_{cf}. The number of distinct fibers of G_{af} and G_{cf} is equal to $|V_{cf}|$ and $|V_{af}|$ respectively.

In addition to acyclic and connectivity properties, a SCOT has two more properties: (i) *Index property*, which emphasizes that the labels of nodes of G_{cf} must be a sequence of unique integers starting from zero, and (ii) *Label Order property*, which requires that the first operand (from left to right) of a $CPUG$ should be G_{af}. The index property implies that the index of each fiber of G_{af} is always an integer. The label order property indicates that the first part of the label of a SCOT node comes from the corresponding vertex of G_{af}, and the second part is the label of the corresponding vertex of G_{cf}, as indicated by Eq. 1. Reversing the order of operands does not generate extra links or nodes, since CPUG is commutative. In Fig. 2, the left operand of $\square$ operator, an acyclic $H - graph$, is the G_{af}, while the second operand G_{cf} is a triangle, which is also a complete graph.

IV. Clustering

In PS systems, clustering is the process of forming groups (i.e., clusters) of brokers dynamically in order to shrink routing paths to reduce delivery delays and IMs [21]. Brokers are clustered based on the similarity of subscriptions of local subscribers, and the frequency of similar messages exchanged among them. Clustering in a SCOT is static, and clusters are identified before an overlay is deployed. Each cluster is assigned a unique identification which is used in the routing of subscriptions and notifications. Each $G^i_{af} - fiber$ in a SCOT is treated as a separate group of brokers called a cluster, and is represented by C_i, where $i \in V_{cf}$. There are always $|V_{cf}|$ number of clusters in a SCOT. The SCOT in Fig. 3 contains three clusters (horizontal layers), each identified with a distinct label C_i, where $i \in \{0,1,2\}$ and is called the *Cluster Index (CI)*. A CI is the label of a vertex of G_{cf} that generates the cluster (or $G_{af} - fiber$) when a CPUG is calculated. Similarly, each $G^j_{cf} - fiber$ is called a *Region* and is represented as R_j, where $j \in G_{af}$ and is known as the *Region Index (RI)*. There are always $|V_{af}|$ number of regions in a SCOT. The SCOT in Fig. 3 contains six regions, each identified by a unique RI.

Overlay Links and Messaging — A SCOT has two types of links: (i) an *intra-cluster overlay link (aCOL)*, and (ii) an *inter-cluster overlay link (iCOL)*. aCOLs connect brokers in the same cluster, while iCOLs connect brokers in the same region. In this paper, solid lines and curves in each cluster in Fig. 3 are aCOLs, connect brokers in each replica of G_{af}. The dotted lines and curves in each region in Fig. 3 are iCOLs, which connect brokers in each replica of G_{cf}. Messaging along aCOLs is referred to as *intra-cluster messaging* while messaging along iCOLs is referred to as *inter-cluster messaging*. L^m_a, the set of all aCOLs in a cluster C_m (where $m \in V_{cf}$) is given as: $L^m_a = \{l\langle (a,m,),(b,m)\rangle | a,b \in V_{af} \wedge ab \in E_{af} \wedge m \in V_{cf}\}$.

The set of all aCOLs in $\mathbb{S}$ is: $\bigcup_{m \in V_{cf}} \mathrm{L}_a^m$. Similarly, L_i^n, the set of all iCOLs in a region R_n for an RI $n \in V_{af}$ is given as: $\mathrm{L}_i^n = \{l\langle(n,u),(n,v)\rangle | n \in V_{af} \wedge u,v \in V_{cf} \wedge uv \in E_{cf}\}$. The set of all iCOLs in $\mathbb{S}$ is expressed as: $\mathrm{L}_i = \bigcup_{j \in V_{af}} \mathrm{L}_i^j$. There are $|E_{af}|.|V_{cf}|$ number of aCOLs and $|V_{af}|.|E_{cf}|$ iCOLs in $\mathbb{S}$, which shows that $E_{\mathbb{S}} = |E_{af}||V_{cf}| + |V_{af}||E_{cf}|$.

Classification of Clusters and Brokers — There are two types of clusters in a SCOT: (i) a *Primary or Host Cluster*, and (ii) a *Secondary Cluster (SC)*. The cluster that contains the host broker of a client is the host or primary cluster of the client, while all other clusters are the secondary clusters of the client. Brokers in a region are direct neighbours, and no two brokers belong to same cluster. This arrangement of brokers requires only one iCOL inter-cluster messaging. The *Primary neighbours* of a broker are its direct neighbours in the same cluster, while the *Secondary neighbours* are those in the same region. Clusters that host interested subscribers of a publisher are called its *Target Clusters (TCs)*. An SC in TCs is called *Secondary Target Cluster (STC)* of the publisher. The host cluster of the publisher is its *Primary Target Cluster (PTC)* if it hosts at least one interested subscriber. An *edge broker* in a SCOT has at most one primary neighbour, while an *inner broker* has at least two primary neighbours. The number of inner or edge brokers in a SCOT is the product of the number of inner or edge brokers in G_{af} with $|V_{cf}|$. All brokers in a region are the same type (i.e., are either inner or edge). A SCOT broker is aware of its own type, the types of its primary and secondary neighbours, and its host CI.

Cluster Bit Vector — Cluster Bit Vector (CBV) is a row (vector) of bits used to identify the host cluster of a subscriber and the STCs of a publisher. CBV is saved with a subscription and is used in dynamic routing of notifications (cf. Section VI). Bits in CBV are indexed from right to left, with the index of the right most bit being zero. Each bit of CBV, called the *Cluster Index Bit (CIB)*, is reserved for a SCOT cluster where the index of the bit is same as the index of the cluster it represents. A subscriber's host broker sets (i.e., turns *ON*) the CIB of the host cluster to 1 before the subscription is forwarded to other brokers. A CBV with a subscription has at most one bit turned ON. For example, the left CBV in Fig. 4 indicates that the subscriber is connected to a broker of cluster C_2. The matching process executed by the publisher's host broker sets CIBs of the STCs to 1 in dynamic routing (cf. Section VI). The CBV of a notification in Fig. 4 (right) indicates that C_0 and C_2 are the STCs. The number of effective bits in a CBV is equal to the number of clusters (i.e., $|V_{cf}|$) in a SCOT. The SCOT in Fig. 3 has only three clusters and thus requires only three bits in CBV to identify STCs. Constructing an underlay–aware cyclic topology for a PS system deployed across different geographic regions is difficult. A SCOT clusters can be deployed in different geographic regions and a cluster's brokers can be placed and connected by respecting some conformity with the local underlay network topology. However, achieving the desired level of conformity is difficult, as all SCOT clusters have the same acyclic topology. A method to relax this restriction for topology structure will be part of the future work.

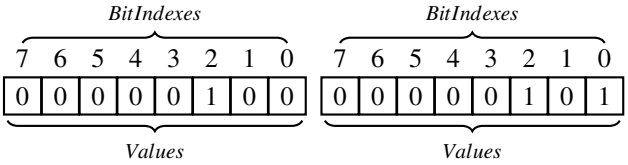

Fig. 4: *(Left) CBV of a subscription indicates that the subscriber is hosted by a broker in cluster C_2. (Right) CBV of a notification indicates that C_0 and C_2 are the secondary target clusters.*

V. Subscription Forwarding

Subscription Forwarding in a SCOT is a two-step process performed by each broker of the host cluster. A subscriber's host broker sets the CIB of the host cluster in CBV to 1. In the first step, a subscription is forwarded to brokers in a subscriber's host cluster. No cycles generate, as the host cluster is a replica of an acyclic factor G_{af} and the subscription forwarding occurs just as in an acyclic overlay. aCOLs of the host cluster are included in a subscription tree generated in the first step of the SFP. In the second step, each broker in the host cluster forwards the subscription to its secondary neighbours, that is, to brokers in the sender's region. The secondary brokers do not forward the subscription any further. In the second step, all iCOLs are added to the subscription tree and the subscription becomes available for all brokers to be matched with notifications. Fig. 5 shows that C_0 is the host cluster of the subscribers S1 and S2, C_1 is the host cluster of S3, and C_2 is the host cluster of S4. The partial subscription trees of S1, S2, and S3, generated in the first step of the SFP, are shown by solid arrow links in the host clusters of the subscribers. The second step of the SFP is illustrated by the dotted arrow links between brokers of the host clusters and their secondary brokers. $l\langle(f,0),(e,0)\rangle$ and $l\langle(b,0),(a,0)\rangle$ are not in the subscription tree of S1. Similarly, $l\langle(a,0),(b,0)\rangle$ and $l\langle(e,0),(f,0)\rangle$ are not in the subscription tree of S2.

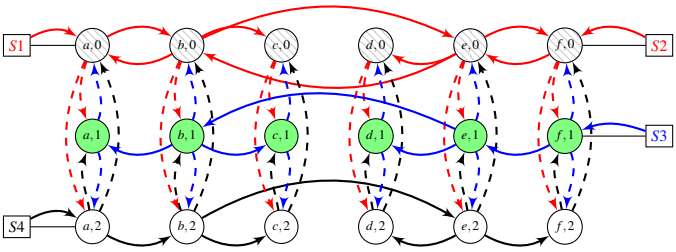

Fig. 5: *The two-step subscription forwarding process in the SCOT shown in Figure 3. The solid arrows indicate part of the subscription tree generated in the first step and the dashed arrows indicate part of the subscription tree generated in the second step.*

The maximum length of a subscription tree generated in the first step is $diam(G_{af})$, while the maximum length of the subscription tree generated after the second step is $diam(G_{af})+1$. In the first step shown in Fig. 5, B(a,0) forwards the subscription of S1 to the primary broker B(b,0), and in the second step, to the secondary brokers B(a,1), and B(a, 2), as indicated by dashed red arrows. Each broker of C_0 repeats the first and the second step for the subscriptions of S1 and S2. The

SFP in SCOT does not generate extra messages to detect and discard duplicates. Each cluster is treated as an exclusive acyclic overlay, and secondary brokers do not forward the subscription. The host cluster of a subscriber is identified through CIB in CBV. Relocation of a subscriber to some other broker in the same host cluster does not generate any updates in RTs of brokers in secondary clusters. This condition ensures that when a broker fails for some reason and its local subscribers relocate to brokers in the same host cluster, no updates propagate to the secondary clusters.

VI. Notification Routing

Notifications are routed to interested subscribers using the RPF technique, following the principle of *downstream replication* [2]. A notification from a publisher may have different delivery delays depending on when the notification is received by the interested subscribers. Three factors, which are important in determining delivery delays are: (i) delay due to in-broker processing to find the next destinations, (ii) length of routing paths, and (iii) notifications payload. We introduce three algorithms for notification routing in a SCOT: (i) a *BID-based* routing algorithm for an unclustered SCOT, (ii) a *Static Notification Routing (SNR)*, and (iii) a *Dynamic Notification Routing (DNR)* algorithms in a cluster-based SCOT.

BID-based Static Routing — An unstructured cyclic *(general)* overlay can have any number of brokers, and each broker can have any number of links. These properties may provide flexibility in the deployment of a PS system that uses an unstructured cyclic overlay. However, the routing algorithms require a notification to carry BIDs assigned to matching subscriptions. [2], [16]. The BID-based approach matches a notification only at the host brokers of a publisher and its interested subscribers, and the notification is forwarded along the paths identified by BIDs (cf. Section II) [16]. We implement a BID-based routing algorithm to compare its performance with the other two routing algorithms. We use BIDs to discard duplicates in the SFP. In order to avoid extra network traffic, subscription trees are generated only once. Links on the routing paths to interested subscribers are known as *target links*. Algorithm 3 in the appendix A provides more information about BID-based routing.

Static Notification Routing — The *Static Notification Routing (SNR)* algorithm uses shortest-lengths subscription trees for routing in a cluster-based SCOT (cf. Section V). SNR algorithm deals with two scenarios: (i) when the host cluster is the only TC of a publisher, and (i) when the publisher has at least one STC. In (i), all interested subscribers are hosted by the brokers in the publisher's host cluster, and lengths of the routing paths satisfy the relation $max\big(d\langle(u_1,v_1),(u_2,v_2)\rangle\big) \leq diam(G_{af})$, where (u_1,v_1) and (u_2,v_2) are any brokers in the publisher's host clusters. In (ii), at least one interested subscriber is hosted by a broker of a secondary cluster, and lengths of the routing paths satisfy the relation $max\big(d\langle(x_1,y_1),(x_2,y_2)\rangle\big) \leq \big(diam(G_{af})+1\big)$, where (x_1,y_1) and (x_2,y_2) are any brokers in a SCOT overlay. In SNR algorithm, the host broker of a publisher forwards notifications to brokers of STCs and to primary neighbours in the routing paths of the interested subscribers in the host cluster. This is further explained in the following.

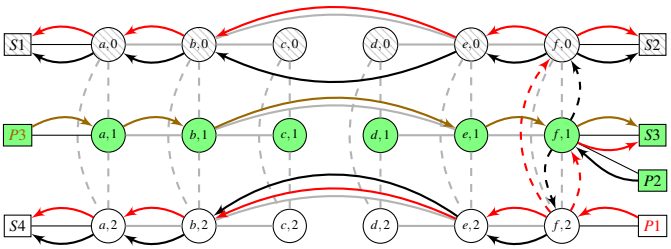

Fig. 6: *Publication routing in SCOT shown in Fig. 3 using the RPF technique. Color of label of a publisher matches with its notification arrows. Solid arrows indicates intra-cluster messaging while dotted arrows represent inter-cluster messaging.*

In Fig. 6, publisher *P1* is hosted by B(f,2) in C_2, while *P2* and *P3* are hosted by B(f,1), B(a,1) in C_1, respectively. The set of interested subscribers corresponding with P1, P2, and P3 are $\{S1, S2, S3, S4\}$, $\{S1, S2, S4\}$ and $\{S3\}$. The set of TCs of P1, P2 and P3 are $\{C_0,C_1,C_2\}$,$\{C_0,C_2\}$ and $\{C_1\}$. A notification from P1 propagates onto the path indicated by the solid red arrows in C_2 to reach S4. B(f,2) creates two additional copies of the notification to forward to brokers (f,1) and (f,0) in the STCs. Solid red arrows in C_1 and C_0 show the notification delivery paths to S1, S2, and S3. Although S1 and S2 are interested in notifications from P1, only one copy of the notification is forwarded to B(f,0) to avoid sending duplicates. Similarly, notifications from P2 and P3 are forwarded to interested subscribers. Note that P2 has no PTC, while P3 has no STC.

Dynamic Notification Routing — *Dynamic Notification Routing (DNR)* refers to the capability of a PS system to alter the routing path in response to overloaded or failed links and/or brokers. DNR requires multiple routing paths from a publisher to interested subscribers. Multiple techniques have been proposed to handle the issue of dynamic routing in address-based networks where the shortest routing paths are calculated from a global view of a network topology graph that is saved on every network router [17]. However, these techniques are not applicable in the PS domain, because brokers are aware of only their direct neighbours and dynamic routing decisions have to be made without having a global view of an overlay. This paper focuses on dynamic routing when one or more iCOLs are overloaded and notifications start queueing up in the output queues of a broker. Overloading and subsequent queuing can happen in two cases: (i) when a broker is not able to process high volume of outgoing notifications and become overwhelm, and (ii) when bandwidth is limited. There is always one subscription-tree per subscription in an SCPS system, which makes DNR difficult. We use the *structuredness* of a cluster-based SCOT to perform inter-cluster DNR for STCs when one or more iCOLs are overloaded. DNR can reduce delivery delays when a high volume of notifications accumulate in the output queues of a broker. We use Eq. 3 to find whether an output queue is congested, and the DNR

should find alternative iCOLs.

$$(Q_\ell).(1+Q_{in}, 1+Q_{out})_{t_w} > \tau \quad (3)$$

Q_{in} is the number of notifications that enter into the output queue, and Q_{out} is the number of notifications that leave the output queue in time window t_w. The term $(1+Q_{in}, 1+Q_{out})_{t_w}$ is the ratio of $(1+Q_{in})$ to $(1+Q_{out})$, and is known as the *Congestion Element (CE)*. CE identifies congestion trend in an output queue in t_w interval. $CE > 1$ indicates that the congestion is increased in the last t_w interval as more messages entered into the output queue than sent out onto the overlay link to next broker. $CE < 1$ shows that congestion is decreased in the last t_w interval, as more messages are sent out onto the overlay link than entered into the output queue. Q_ℓ is the length of the output queue (i.e., the number of notifications waiting in the output queue to be forwarded to their next destinations). An output queue is congestion–free when CE is 1 and Q_ℓ is 0. Octopi saves the values of Q_{in} and Q_{out} in a *Link Status Table (LST)* on each broker, and the values are updated after each t_w interval. DNR in a SCPS system becomes active when Eg. 3 is valid for some positive real value of τ in t_w interval. In the SNR algorithm, a publisher's host broker adds exclusive copies of a notification in the output queues of the target links. For example, if there are α target aCOLs and β target iCOLs for a publisher, which generates γ number of notifications in t_w interval, then the publisher's host broker adds an $(\alpha+\beta).\gamma$ number of notifications in the output queues of the target links in t_w interval. A High Rate Publisher (HRP) with a very high value of γ can overwhelm the host broker and the brokers in routing paths of interested subscribers. DNR algorithm can alleviate overwhelmed brokers by adding fewer copies of a notification in the congested output queues of the target iCOLs. In particular, DNR adds *no* copy of a notification in the congested output queues when at least one output queue of a target aCOL or iCOL is uncongested. Furthermore, when the output queues of all target links are congested, DNR adds only one copy of the notification to the least congested output queue of a target iCOL. The mechanism of DNR is further explained by using three different cases. Fig. 7 shows a SCOT whose G_{af} is a *path graph* of three nodes and G_{cf} is a triangle. The set of subscribers $\{S1, S2, S3, S4\}$ is interested in notifications from a publisher P.

(i) When the output queues of all target iCOLs are congested and at least one output queue of a target aCOL is uncongested, DNR sets the CIBs of the STCs to 1 and adds the CBV in the header of a notification. The notification is forwarded onto aCOL of the uncongested output queue. Using this technique, DNR algorithm keeps the length of the congested output queues of the target iCOLs short. The load of forwarding the notification to STCs is shifted to the next primary broker using the heuristic that unoverloaded target iCOLs are available down the routing path. The number of notifications added to the output queues (of the notification sending broker) of the target links in t_w interval is $\alpha\gamma$. Fig. 7(a) indicates that the two iCOLs, $l\langle(b,2),(b,1)\rangle$ and $l\langle(b,2),(b,0)\rangle$, are overloaded (thick links), and a notification is sent only to S3 and S4 in the host cluster. Since two unoverloaded aCOLs are available, the aCOL with the least Q_ℓ is selected to add CBV with bits 011 (presumably, the aCOL $l\langle(b,2),(c,2)\rangle$ has least value of Q_ℓ). This approach attempts to maintain a balance between lengths of the output queues. The notification which carries CBV is called *the CBV-Notification (or CBV-N)*, and is distinguished by the red arrow messages in Fig. 7. As the target iCOLs are not overloaded at B(c,2), the notification is routed to the STCs. To forward one notification from P in this case, the DNR algorithm generates 4 IMs, while SNR generates 6 IMs. Furthermore, overloaded iCOLs $l\langle(b,2),(b,0)\rangle$ and $l\langle(b,2),(b,1)\rangle$ are excluded from the routing paths, which is not possible with the SNR algorithm.

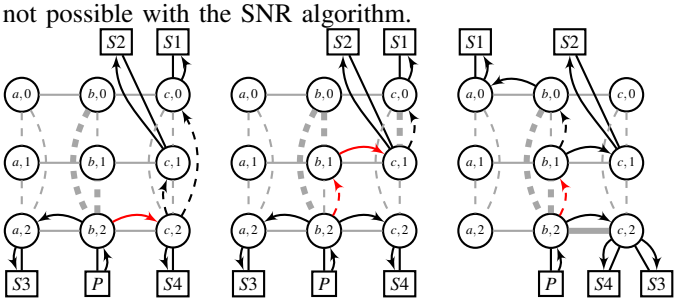

(a) DNR Case (i) (b) DNR Case (ii) (c) DNR Case (iii)

Fig. 7: *Dynamic routing in SCOT with 3 clusters and 9 brokers. Grey solid lines inicate aCOLs while dotted lines and curves represented iCOLs. Thick aCOL and iCOLs represent overloaded links. Red arrow messages represent CBV-N.*

(ii) CBV-N is forwarded onto the unoverloaded target iCOL. If more than one unoverloaded target iCOL is available, CBV-N is forwarded onto the target iCOL with the least value of Q_ℓ. The number of notifications added to the output queues of the target links in t_w interval is $(\alpha+\theta)\gamma$, where θ is the number of unoverloaded target iCOLs and $\theta < \beta$. Fig. 7(b) indicates that the CBV-N, with CBV 001, is forwarded to $l\langle(b,2),(b,1)\rangle$). Since the link $l\langle(b,1),(b,0)\rangle$ is overloaded, the CBV-N is forwarded to B(c,1) using the subscription tree of S2. The link $l\langle(c,1),(c,0)\rangle$ is also overloaded but no aCOL is available to search an unoverloaded iCOL, and the notification is forwarded to B(c,0) onto the overloaded iCOL $l\langle(c,1),(c,0)\rangle$. To forward one notification from P to interested subscribers, the DNR algorithm generates 5 IMs, while the SNR algorithm generates 6 IMs. Although DNR excludes the overloaded links $l\langle(b,2),(b,0)\rangle$ and $l\langle(b,1),(b,0)\rangle$ from the routing path, the notification has to be forwarded to C_0 onto the overloaded link $l\langle(c,1),(c,0)\rangle$, and the dynamic routing path generated to send the notification to S1 is processed by one additional broker.

(iii) If all target links are overloaded, the CBV-N is forwarded onto the least loaded target iCOL. The number of notifications added to the output queues of the target links in t_w interval is $(\alpha+1)\gamma$. Fig. 7(c) shows that CBV-N is forwarded onto the least overloaded link $l\langle(b,2),(b,1)\rangle$ with CBV is 001. The number of IMs generated by SNR and DNR algorithms is 5; however, the dynamic route used to send the notification to S1 has one additional broker. The principle of downstream replication is followed when the notification is forwarded onto

the link $l\langle(b,2),(c,2)\rangle$. DNR is a best-effort algorithm for inter-cluster dynamic routing and depends on the subscription trees laid-on by interested subscribers, which are hosted by the routing cluster. The algorithm does not guarantee finding an unoverloaded iCOL, even if one exists. In Fig. 7(c), DNR does not use the unoverloaded link $l\langle(b,2),(a,2)\rangle$ for dynamic routing because the link is not a target aCOL. DNR also does not support *intra-cluster dynamic routing*. An algorithm which covers both these cases will be presented in an extended version of this paper.

VII. Evaluation

We implemented the BID-based routing, SNR, and DNR algorithms in Octopi. Octopi is developed on top of PADRES, an open source advertisement-based CPS system [7]. PADRES messages are serialized Java objects, which consists of two parts: the message header and the message body. Octopi is an SCPS variant of PADRES and generates only subscription trees for notification routing. Octopi uses the language model and matching logic of PADRES, but features like patterns and sequences of events are not supported. PADRES offers content-based routing in acyclic and cyclic (general) overlays, however, Octopi currently works with structured cyclic overlays (i.e., SCOT).

Setup — For simplicity, we discuss only the factors of $\mathbb{S}$ (the topology we used for experimental evaluation) which are shown in Fig. 9. The G_{af} factor of $\mathbb{S}$ is an acyclic topology of 15 brokers and the G_{cf} factor of $\mathbb{S}$ has 5 brokers. Furthermore, the G_{af} factor has 5 inner brokers (*vi, vii, viii, ix* and *x*) and 10 edge brokers. Since each of these brokers has $|V_{cf}|-1$ number of secondary neighbours, there are 25 inner brokers and 50 edge brokers (for a total of 75 brokers) in $\mathbb{S}$. Brokers of $\mathbb{S}$ were deployed on a cluster of 10 computing nodes, where each computing node consisted of 2 Intel(R) Xeon(TM) E5-2620 CPU with total 12 cores of 2.1 GHz each, and 64 GB RAM. One Mellanox SX1012 switch with 10 Gbps links were used to connect the computing nodes. $\mathbb{S}$ was deployed in such a way that the primary and the secondary neighbours of each broker were always on different computing nodes. Each broker was loaded in a separate instance of JVM with 1 GB initial memory.

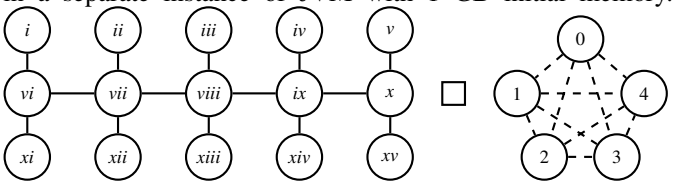

Fig. 9: *Left operand of □ operator is the G_{af} factor while the right operand is the G_{cf} factor of $\mathbb{S}$. The G_{cf} generates five replicas of G_{af}.*

Stock datasets are commonly used to generate workloads for evaluations of CPS systems [8]. We used a dataset of 500 stock symbols from Yahoo Finance!, and each stock notification had 10 distinct attributes. This high dimension data require high computation for filtering information during in-broker processing. Subscriptions were generated synthetically. We randomly distributed publishers and subscribers to brokers in $\mathbb{S}$, and each subscriber registered only one subscription.

Metrics — Through experiments with real world data, we evaluated Octopi using primitive metrics, such as subscription, notification, and matching delays.

Subscription Delay: The subscription (forwarding) delay is the maximum time elapsed as a subscription message reach any broker in an overlay. A subscription is expected to take less time to reach a broker in close proximity to the host broker of a subscriber than to reach a broker in a far-off region of the overlay. Since the SFP in a cluster-based SCOT generates shortest-length subscription trees, it is important to measure the difference between the average subscription delays in unclustered and cluster-based SCOT.

Notification Delay: The notification delay measures end-to-end latency from the time a notification is generated to the time it is received by a subscriber. As the average length of subscription trees in the BID-based approach is higher than the average length of the subscription trees in a cluster-based SCOT approach, knowing the difference in latencies in these two cases is a worthwhile area of inquiry.

Matching Delay: The matching delay is the time taken by each broker in a routing path to find subscriptions that match with the notification contents. A notification in BID-based routing is matched at the host brokers of a publisher and its interested subscribers.

Inter-broker Messages (IMs): The number of IMs depends on the lengths of subscription trees that the SFP generates, as well as the relative distance between publishers and subscribers. As the average length of subscription trees in cluster-based SCOT is expected to be less than the average length in unclustered SCOT, the number of IMs generated by the SNR and the DNR algorithms is expected to be less than the number of IMs generated by the BID-based routing algorithm.

Results — The results presented in this section cover three important aspects of evaluation: (i) Subscriber Scalability, which studies the behaviour of the routing algorithms when the number of subscribers increases while the number of publishers remains constant; (ii) Publisher Scalability, which is connected with the study of the algorithms with a varying number of publishers and a constant number of subscribers; (iii) Burst Scenario, in which an HRP starts sending notifications at a high rate and causes congestion in the output queues. Publishers and subscribers were randomly distributed among brokers of $\mathbb{S}$.

Subscriber Scalability: We increased the number of subscribers from 500 to 5000, and each subscriber registered one subscription. We used 100 publishers, each sending 1000 notifications at the rate of 60 notifications per minute. All publishers started sending notifications in the first 5 seconds after all the subscriptions registered. Fig. 8(a) shows the average subscription delay in unclustered $\mathbb{S}$ (legend BID) is higher than the cluster-based $\mathbb{S}$ (legend SNR and DNR). Importantly, there is no difference in the subscription delays when SNR and DNR are used, since the subscription forwarding algorithm does not use DNR in the SFP, even if some output queues are congested. Fig. 8(a) indicates that the average subscription

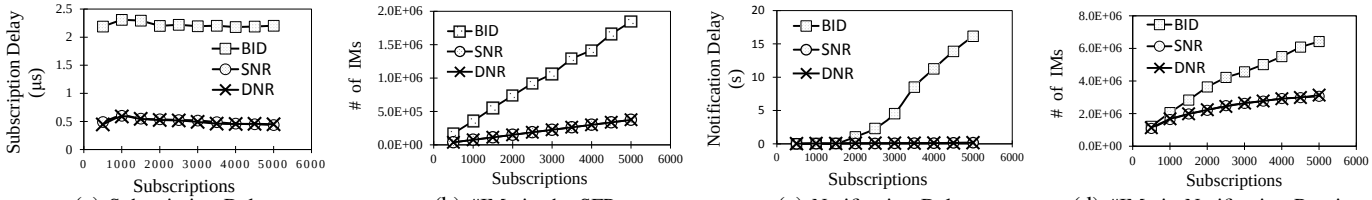

(a) Subscription Delay (b) #IMs in the SFP (c) Notification Delay (d) #IMs in Notification Routing

Fig. 8: *Subscriber Scalability in an unclustered and cluster-based SCOT*

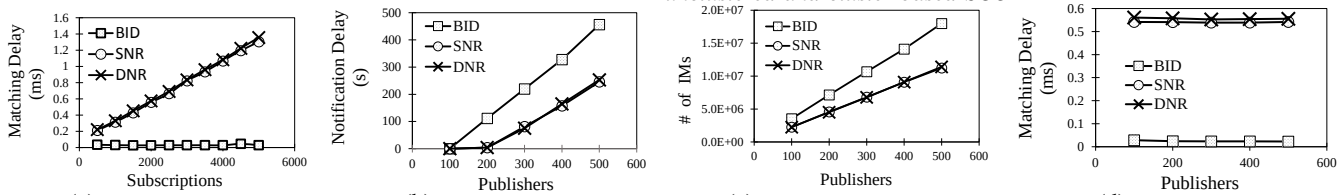

(a) Matching Delay (b) Publication Delay (c) Inter-Broker Messages (d) Matching Delay

Fig. 10: *(a) Matching dealy in Subscriber Scalability. (b),(c),(d) Publisher Scalability in an unclustered and cluster-based SCOT.*

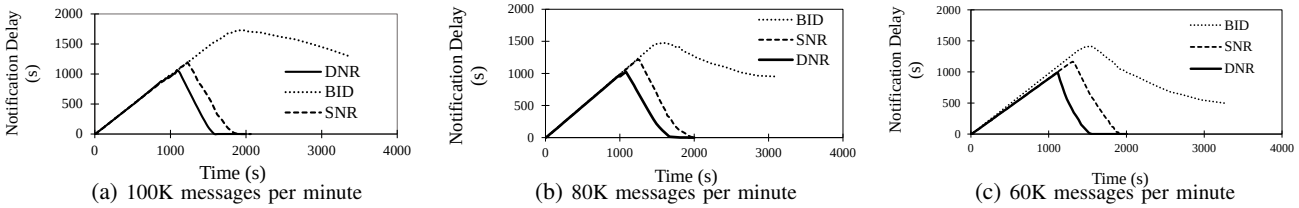

(a) 100K messages per minute (b) 80K messages per minute (c) 60K messages per minute

Fig. 11: *Stability analysis.*

delay in unclustered and cluster-based $\mathbb{S}$ is nearly constant. However, the average subscription delay in an unclustered $\mathbb{S}$ is 77% higher. This difference is mainly due to the larger lengths of subscription trees generated by the SFP in unclustered $\mathbb{S}$. Fig. 8(b) shows that the number of IMs generated in the SFP in unclustered $\mathbb{S}$ is higher than cluster-based $\mathbb{S}$. More specifically, the SFP generates 4.86 times more IMs in unclustered $\mathbb{S}$. There are two reasons for this significant difference: (i) higher average lengths of subscription trees in unclustered $\mathbb{S}$, and (ii) extra IMs generated to discard duplicate subscriptions (cf. Section II). Fig. 8(c) shows that the growth in end-to-end notification delay in a BID-based approach is exponential as compared to a very small increase in SNR and DNR algorithms. More specifically, the average notification delay in BID-based algorithm is 59 times higher than the average delay in SNR and DNR algorithms. There are two reasons for this large difference: (i) the average length of subscription trees in BID-based approach is higher than cluster-based algorithms, and (ii) increase in payload due to carrying BIDs in notifications. More BIDs may have to be added to notifications with increased number of subscribers (and brokers). Fig. 8(d) shows that the number of IMs generated to send the same number of notifications (1000 per publisher) in the BID-based approach is 72.7% more than in SNR and DNR. Again, the contrast is due to the difference in the lengths of subscription trees. Fig. 10(a) shows that the average matching delay in SNR and DNR algorithms increases linearly, while it is nearly constant in BID-based routing algorithm. More specifically, the average matching delay in SNR and DNR algorithms is 26.77 times higher than in BID-based routing algorithm. In BID-based notification routing, only the host brokers of a publisher and its interested subscribers execute the matching process to find BIDs of brokers and matching subscriptions which should receive a copy of a notification. No matching occurs at the intermediate brokers. However, in the SNR and DNR algorithms, matching is performed at each broker in the routing path.

Publisher Scalability: We increased the number of publishers from 100 to 500, each sending 500 notifications at a fixed rate of 100 notifications per minute. The number of subscribers in this experiment was 2000. Fig. 10(b) shows that the average notification delay in BID-based routing is higher than in SNR and DNR routing. More specifically, the average delay in the SNR and DNR algorithms is 61% less than the BID-based approach. The difference is due to the larger length of subscription trees and higher payload from carrying BIDs in notifications. Due to the difference in lengths of subscription trees, the number of IMs generated in BID-based routing is 36.4% higher than in the SNR and DNR algorithms (Fig. 10(c)). The average matching delay in SND and DNR algorithms is 22.5 times higher than in BID-based algorithm (Fig. 10(d)).

Stability Analysis: The stability analysis tells how quickly a CPS system converge to a normal state after an HRP finishes sending notifications. To study this behaviour, we set the value of τ to 10 and t_w to 50 milliseconds. We used 2000 subscribers and 100 publishers in this experiment. Each subscriber had 2% selectivity, while each publisher issued 2000 notifications at the rate of 60 per minute rate. 2% of the subscribers subscribed to receive notifications from an HRP. We executed three simulations in which the HRP sent 100K notifications at rates of 100K, 80K, and 60K notifications per minute. The HRP and

its interested subscribers did not connect to the same broker and all clusters of $\mathbb{S}$ were TCs of the HRP. This arrangement makes sure that each cluster receives notifications from the HRP. Furthermore, the HRP and its interested subscribers were hosted by different brokers exerting more load on iCOLs and aCOLs. The burst of notifications started after all the subscriptions were registered in $\mathbb{S}$. The burst continued from 60 to 100 seconds depending on the rate of the HRP. Each point in the graphs in Fig. 11 is a maximum delivery delay of 1000 notifications received in a sequence. This approach helps in analysing the stability of a SCPS system without graphing too many points. Each simulation was run until all notifications (6 million in total) were received by subscribers. Fig. 11 shows that DNR stabilized Octopi before the SNR algorithm, while the BID-based routing algorithm was not able to stabilize the system for the same workload. On average, in the first 18 minutes and 30 seconds, the maximum delay of the notification (out of 1000) was the same in the three routing algorithms. No tendency toward stability was observed. This indicates that, due to the high rates of notifications from the HRP, the state of the system (Octopi) did not converge to normal until the condition $CE < 1$ was maintained for some time (on average, 16 minutes and 40 seconds for the three simulations). DNR started stabilizing Octopi before other two algorithms. The average value of Q_ℓ of target links at the host broker of HRP when DNR was used was 61% less than SNR and 73% less than BID-based routing algorithms. There were 5 TCs of the HRP and DNR tends to add the minimum number of copies of a notification when the output queues of the target links are congested. The average notification delay in the DNR algorithm when the notification rate was 100K was 17.5% less than the SNR algorithm. Similar improvements, when the rates were 80K, 60K messages per minute, were 17.7%, and 21%, respectively. On average, DNR algorithm stabilized the system 239 seconds before the SNR algorithm and generated only 0.3% IMs more than the SNR algorithm and 58.6% less than the BID-based routing algorithm. Analysis of the collected data indicates that the number of messages that had delivery delays less than 1 second in the DNR, SNR and BID-based routing algorithms are 53.8%, 44.1% and 31%, respectively. We also conducted several experiments with HRP on each cluster and sending notifications with different rates. The performance difference between DNR and SNR decreased with an increase in the number of HRPs and decrease in their burst rates.

VIII. Conclusion

This paper introduces the first subscription-based CPS system that offers inter-cluster dynamic routing of notifications. CPUG is used to define a structured cyclic overlay SCOT to generate subscription trees of the shortest lengths and eliminate cycles in notification routing. The main contribution of this paper is the DNR algorithm that uses a lightweight bit vector mechanism with cluster-based SCOT to provide dynamic routing of notifications between host clusters of interested subscribers. The algorithm does not require a global knowledge of an overlay topology and dynamic routing is done when congestion is detected in output queues of a broker. An implementation of the DNR and SNR algorithms in PADRES indicates that the DNR and SNR scale well with the number of publishers and subscribers.

Acknowledgements — We gratefully acknowledge the support of the Canada Foundation for Innovation in providing the computing equipment used for this work. Moreover, we would like to thank Martin Karsten, Naomi Nishimura and others for their feedback on our initial work.

References

[1] César Cañas, Kaiwen Zhang, Bettina Kemme, Jörg Kienzle, and Hans-Arno Jacobsen. Publish/subscribe network designs for multiplayer games. In *Proceedings of the 15th International Middleware Conference*, Middleware '14, pages 241–252, New York, NY, USA, 2014. ACM.

[2] Antonio Carzaniga. Architectures for an event notification service scalable to wide-area networks, phd thesis, politecnico di milano. 1998.

[3] Alex King Yeung Cheung and Hans-Arno Jacobsen. Load balancing content-based publish/subscribe systems. *ACM Trans. Comput. Syst.*, 28(4):9:1–9:55, December 2010.

[4] Yogen K. Dalal and Robert M. Metcalfe. Reverse path forwarding of broadcast packets. *Commun. ACM*, 21(12):1040–1048, December 1978.

[5] K. Efe and A. Fernandez. Products of networks with logarithmic diameter and fixed degree. *Parallel and Distributed Systems, IEEE Transactions on*, 6(9):963–975, Sep 1995.

[6] Cooper et al. Pnuts: Yahoo!'s hosted data serving platform. *Proc. VLDB Endow.*, 1(2):1277–1288, August 2008.

[7] Hans-Arno et al. The PADRES Publish/Subscribe System. In *Principles and Applications of Distributed Event-Based Systems*, pages 164–205. IGI Global, 2010.

[8] Navneet et al. Minimizing the Communication Cost of Aggregation in Publish/Subscribe Systems. In *35th IEEE International Conference on Distributed Computing Systems (ICDCS)*, 2015.

[9] Yogeshwer Sharma et al. Wormhole: Reliable pub-sub to support geo-replicated internet services. In *NSDI 15*, pages 351–366, Oakland, CA, May 2015. USENIX Association.

[10] Patrick Th. Eugster, Pascal A. Felber, Rachid Guerraoui, and Anne-Marie Kermarrec. The many faces of publish/subscribe. *ACM Comput. Surv.*, 35(2):114–131, June 2003.

[11] Richard Hammack, Wilfried Imrich, and Sandi Klavzar. *Handbook of Product Graphs, Second Edition*. CRC Press, Inc., Boca Raton, FL, USA, 2nd edition, 2011.

[12] Microsoft Bing Pulse (http://pulse.bing.com/).

[13] Google Cloud Pub/Sub (https://cloud.google.com/pubsub/).

[14] Reza Sherafat Kazemzadeh. *Overlay Neighborhoods for Distributed Publish/Subscribe Systems*. PhD thesis, University of Toronto, 2012.

[15] Jay Kreps, Neha Narkhed, and Jun Rao. Kafka: a distributed messaging system for log processin. *NetDB Workshop*, pages 01–07, 2011.

[16] Guoli Li, Vinod Muthusamy, and Hans-Arno Jacobsen. Adaptive content-based routing in general overlay topologies. Middleware '08, pages 1–21, New York, NY, USA, 2008. Springer-Verlag New York, Inc.

[17] Deepankar Medhi and Karthikeyan Ramasamy. *Network Routing: Algorithms, Protocols, and Architectures*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2007.

[18] Vinod Muthusamy. *Flexible Distributed Business Process Management*. PhD thesis, University of Toronto, 2011.

[19] E. Parsonage, H.X. Nguyen, R. Bowden, S. Knight, N. Falkner, and M. Roughan. Generalized graph products for network design and analysis. In *Network Protocols (ICNP), 2011 19th IEEE International Conference on*, pages 79–88, Oct 2011.

[20] Muhammad Adnan Tariq, Boris Koldehofe, Sukanya Bhowmik, and Kurt Rothermel. Pleroma: A sdn-based high performance publish/subscribe middleware. Middleware '14, pages 217–228, New York, NY, USA, 2014. ACM.

[21] Sasu Tarkoma. *Publish / Subscribe Systems: Design and Principles*. Wiley Publishing, 1st edition, 2012.

[22] Kaiwen Zhang, M. Sadoghi, V. Muthusamy, and H.-A. Jacobsen. Distributed ranked data dissemination in social networks. In *Distributed Computing Systems (ICDCS), 2013 IEEE 33rd International Conference on*, pages 369–379, July 2013.

IX. Appendix A

We wrote 6500 lines of Java code to implement Octopi on top of PADRES. The list of additional tools that we developed for this research includes: a SCOT topology generator, which requires topology description in a text input file; a data analysis utility tool, which efficiently analyses high volume of log data. This section further provides the high level details of the algorithms.

Subscription Forwarding in Unclustered SCOT — An unclustered $\mathbb{S}$ is used to evaluate BID-based routing. The SFP in unclustered SCOT is performed using the Algorithm 1. The *unclusteredScotSFP(s)* method saves a subscription in RT and forwards it to adjacent brokers (direct neighbours) if the subscription is received for first time (line 2). Only host brokers assign BIDs to their local subscriptions (lines 3-5).

Algorithm 1: $unclusteredScotSFP(s)$

```
   Input: s : a subscription message;
1  return DL : a list of next destinations that should receive s;
2  if s.bid ∉ RT ∧ s ∉ RT then
3      if isHostBroker(s) then
4          s.bid ← getBID();
5          s.hostBroker ← false;
6      RT.add(s);
7      foreach n ∈ (NeighbourList − s.sender) do
8          s.next ← n;
9          DL.add(s);
```

Subscription Forwarding in Cluster-based SCOT — Subscription forwarding in a cluster-based SCOT is a two step process. In the first step, each broker of the subscriber's host cluster forwards the subscription to its primary neighbours (excluding the one which sends the subscription). As the topology of a SCOT cluster is acyclic, no cycles and duplicates appear in this step. In the second step, each broker of the host cluster forwards the subscription to secondary neighbours in the same region. Secondary brokers do not forward the subscription any further. Algorithm 2 shows that a subscription message has two states (saved in *state* attribute): *PRIMARY* and *SECONDARY*. The state of the subscription message is *PRIMARY* in the first step of the SFP. The host broker of a subscriber turns the CIB of the host cluster *ON*, saves the CBV in the subscription message header (lines 2-5), and forwards the subscription to direct neighbours. The *isPrimary(n)* method is used to check the type of the next broker and the state attribute of the subscription is set accordingly (lines 5, 8-9). The subscription message is saved on the primary and secondary brokers (lines 6,13). The secondary brokers save the subscription, with the state attribute set to *SECONDARY*, and do not forward it to any other broker.

Algorithm 2: $scotSFP(s)$

```
   Input: s : a subscription message;
   Output: DL : a list of next destinations that should receive s ;
1  if s.state = PRIMARY ∨ isHostBroker(s) then
2      if isHostBroker(s) then
3          s.CBV ← setBitOn(CI);
4          s.hostBroker ← false;
5          s.state ← PRIMARY;
6      RT.add(s);
7      foreach n ∈ (NeighbourList − Sender) do
8          if isPrimary(n) ≠ true then
9              s.state ← SECONDARY;
10         s.next ← n;
11         DL.add(s);
12 else if s.state = SECONDARY then
13     RT.add(s) ;
```

BID-based Notification Routing — Algorithm 3 provides a snapshot of the BID-based notification routing. For routing of each notification from a local publisher, only host broker of the publisher embeds the BIDs of the matching subscriptions (lines 2-5). BIDs of the matching subscriptions are saved in *bidList* attribute of a notification p. When received by any broker in the routing path, the *splitSTIDs(bidList)* method creates hash map which uses next–destination path as hash key and the list of BIDs of brokers down to the next destination as object of the hash map. The presence of an object *BID* in the hash map indicates that some local subscribers have matching subscriptions and should receive a copy of the notification (lines 7-12). Afterwards, a copy of the notification with the corresponding list of the BIDs are forwarded to all next destinations (lines 13-15).

Static Notification Routing (SNR) — The host broker of a publisher forwards notifications from the publisher to interested subscribers in target clusters. Algorithm 4 provides details about static notification routing in $\mathbb{S}$. The method *getDistinctSubs(p)* returns a list of distinct matching subscriptions *(IS)* (line 3) based on their last hop to avoid sending duplicates. Like a subscription message, the publisher's host

Algorithm 3: $pubBID(p)$

```
   Input: p : a notification message;
1  return DL : a list of next destinations that should receive p ;
2  if isHostBroker(p) then
3      IS ← getInterestedSubs(p);
4      foreach s ∈ IS do
5          p.bidList ← s.bid;
6  nextBIDs ← splitBIDs(p.bidList);
7  if BID ∈ nextBIDs.objects then
8      localSubs ← getLocalSubs(p);
9      foreach client ∈ localSubs do
10         p.next ← client;
11         DL.add(p);
12     nextBIDs ← nextBIDs − DIB;
13 foreach nextHop ∈ nextBIDs.keys do
14     p.next ← nextHop;
15     p.bidList ← nextBIDs.get(nextHop);
16     DL.add(p);
```

broker assigns two states to a notification. A notification has *PRIMARY* state in the primary cluster of its publisher and *SECONDARY* state in the STCs. The state of a notification is set using CI of the host cluster of an interested subscriber (lines 5-8). The method *nextLocalDestinations* provides the list of the next destinations in the same cluster (either the primary cluster or the secondary cluster). Only one copy of the notification is forwarded onto each link to avoid duplicates (lines 12-15).

Algorithm 4: $scotSNR(p)$

```
   Input: p : a publication message;
1  return DL : a list of next destinations that should receive p;
2  if isHostBroker(p) then
3      IS ← getDistinctSubs(p);
4      foreach s ∈ IS do
5          if s.CIB ≠ CI then
6              p.state = SECONDARY;
7          else
8              p.state = PRIMARY;
9          p.next ← s.lastHop;
10         DL.add(p)
11 else
12     destinations ← nextLocalDestinations(p);
13     foreach d ∈ destinations do
14         p.next ← d;
15         DL.add(p)
```

Dynamic Notification Routing (DNR) — Upon receiving a notification, each broker prepares a list of local subscribers (line 7), and a list of direct neighbours (next destinations), represented by $\mu 1$, which should receive the notification. Both lists include only unique elements to avoid duplicates. $\eta 1$ represents the list of the overloaded iCOLs to the next destinations and $\eta 2$ is least load iCOL to any of the next destinations. The $getLeastLoaded_aCOL(\mu 1)$ returns the least loaded link ℓ among the next aCOL destinations. *local* represents the list of local subscribers, which receive a notification at the beginning of forwarding process (line 9-11). Next, the CIBs of the overloaded links, which are represented by η, are set to 1 in the CBV (lines 12-15). Note that the CIB of the notification routing in local cluster is not set to 1 (line 14), as the DNR does not support *intra-cluster dynamic routing*. $\mu 2$ is the list of unoverloaded direct neighbours that should receive a copy of the notification (16-18). Note that the condition $(\mu 2 - \eta 1)$ on the loop prevents sending the notification onto overloaded iCOLs. As in the last part of the algorithm, the CBV-N is identified using the most suitable link (cf. VI) (lines 19-30).

Algorithm 5: $scotDNR(p)$

```
   Input: p : a publication message;
1  return DL : a list of next destinations that should receive p;
2  IS ← getInterestedSubs(p);
3  μ1 ← nextUniqueDestinations(IS);
4  η1 ← getOverloaded_iCOLs(μ1);
5  η2 ← getLeastLoaded_iCOL(μ1);
6  ℓ ← getLeastLoaded_aCOL(μ1);
7  local ← getLocalSubscribers(μ1);
8  CBV ← 0;
9  foreach d ∈ local do
10     p.next ← d;
11     DL.add(p);
12 μ2 ← (μ1 − local);
13 foreach d ∈ η do
14     if d.CIB ≠ CI then
15         CBV ← setBitOn(CBV, d.CIB);
16 foreach d ∈ (μ2 − η1) do
17     p.next ← d;
18     DL.add(p);
19 if CBV ≠ 0 then
20     if isOverloaded(η2) = false then
21         p.next ← η2;
22         DL.add(p);
23     else if isOverloaded(ℓ) = true then
24         ℓ ← getLeastLoaded_iCOL(η);
25         p.next ← ℓ;
26         DL.add(p);
27     else
28         foreach msg ∈ DL do
29             if msg.next = ℓ then
30                 msg.CBV ← CBV;
```